

PATENT APPLICATION

I/O METHOD AND APPARATUS FOR OPTICAL STORAGE MEDIA

Inventors:

Pontus Söderström
Karl XI gatan 8a
222 20 Lund
046-211 56 35
A Citizen of Sweden

Per Sten
Linnég. 9
224 60 Lund
046-127479
A Citizen of Sweden

Thu Le CHI
Jägaregatan 41
226 52 LUND
046-151124
A Citizen of Sweden

Assignee:

AXIS AB
Scheelevägen 34
SE-223 63 Lund, Sweden
A Swedish Corporation

CARY & KELLY, LLP
1875 Charleston Road
Mountain View, CA 94043
Telephone (650) 316-4009

003021" / 0455/60

One disadvantage of CD-ROM drive technology is that the access times and transfer rates associated with CD-ROM drives are in general substantially slower than other comparable data storage technologies. An ordinary double speed CD drive commonly found in many of today's computers has an access time of 300 milliseconds. In contrast a comparable hard disk commonly found in many of today's computers has an access time of 10 milliseconds. Furthermore ordinary dynamic read-only memory (DRAM) found in many of today's computers has an access time of 60 ns. The amount of time required to read and write data to the CD drive is much greater than a corresponding time required to access other data storage technologies. With a large speed discrepancies between optical storage medium and other types of data storage there is a need to improve the performance, scalability, and availability of CD-ROM drives.

SUMMARY OF THE INVENTION

The present invention provides a network enabled multi-user configurable optical storage server. The device is easy to install and does not require client or server software. The server can be configured across a network through a series of HTML based web pages, or graphical user interfaces (GUI) provided by the server.

The server can be accessed across a network by one or more users, terminals, or other network attached devices for substantially concurrent read/write access to one or more optical storage media. The server supports a comprehensive set of networking protocols and operating systems simultaneously. The server may support from one to hundreds of CD/DVD-ROMS.

In an embodiment of the invention an apparatus for providing I/O access to at least one data storage device across a network is disclosed. The apparatus includes a network module, a first and second stage and a logic. The network module is configured to couple to the network for sending and receiving data packets. The first and second stages serially couple the network module to the storage device. The logic is responsive to a received data packet from the network to serially move data from the received data packet from said network module through each of said first and said second stages to the at least one storage device. The logic is further responsive to a read request from the network to send a data packet to the network via said network module from whichever of said at least first and second stages and the at least one storage device includes the data.

In an alternate embodiment of the invention an apparatus for providing I/O access to at least one data storage medium across a network is disclosed. The apparatus includes a network module, at least one data storage device a hard drive and logic. The network module is configured to couple to the network for sending and receiving data packets. The at least one data storage device is coupled to the network module and the at least one data storage device for providing an input and an output of datum stored on the at least one data storage medium. The hard drive is coupled to the network module for the caching of at least one of file structures for the stored datum and file structures together with the corresponding datum stored on the at least one data storage medium. The logic is for determining a selected cache policy for the at least one data storage device based on a user selection and for caching on said hard drive a corresponding selected one of the file structures and file structures together with the corresponding datum responsive to the user input.

In an embodiment of the invention a method for providing I/O access to at least one data storage device across a network is disclosed. The method comprises the acts of:

coupling to the network for input and output of data;

coalescing data received from the network which corresponds with a selected file or data stream;

008021 120800 09733707

storing said data aggregated in said act of aggregating.

coupling to the network for sending and receiving data packets;

10 caching at least one of file structures for the stored datum and file structures
together with the corresponding datum stored on the at least one data storage medium
responsive to a user selection.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, wherein like reference numerals
5 designate like structural elements, and in which:

FIG. 1 shows a plurality of devices coupled across the network with an optical storage server.

FIG. 2 is a detailed hardware block diagram of the server shown in FIG. 1 including graphical user interfaces stored therein for server setup over the network.

10 FIGS. 3A-F are combined hardware block diagrams of the server and optical storage media shown in FIG. 1 during a write phase of I/O operation.

FIGS. 4A-B are process flow diagrams corresponding with the writing and reading of data to an optical or DVD media as shown in FIGS. 3A-F.

15 FIG. 5 is a hardware block diagram of the server, optical and electronic storage mediums shown in FIG. 1.

FIG. 6A-C show the data structures associated with both the CD and hard drives shown in FIG. 5.

20 FIGS. 7AB are process flow diagrams corresponding with the boot phase and operational phase of operation of the server shown in FIGS. 5-6AC.

DETAILED DESCRIPTION OF THE INVENTION

The present invention provides a network enabled multi-user configurable optical storage server. The device is easy to install and does not require client or server software. The server can be configured across a network through a series of
25 HTML based web pages, or graphical user interfaces (GUI) provided by the server. The server can be accessed across a network by one or more users, terminals, or other network attached devices for substantially concurrent read/write access to one or more optical storage media. The server supports a comprehensive set of networking protocols and operating systems simultaneously. The server may support from one to

09733707-120800

hundreds of CD/DVD-ROMS. In an embodiment of the invention the server includes two stages of writing prior to committing to a write to the physical medium. This speeds both read and write access to the disk, allows concurrent user access to the disk, and greatly simplifies the mapping and file structures on the disk. In an alternate
 5 embodiment of the invention the server includes a hard drive for configurable caching of selected ones of the optical storage media.

FIG. 1 shows a plurality of devices coupled across the network with an optical storage server. Computers for 118, 120 and scanners 124,134 interface across the network 100 with the optical storage server 102. The optical storage server includes a
 10 server portion 106 and one or more optical storage drives, e.g. CD drives 108,110, and 112. Optical disks 104 may be inserted and removed from any or all of the above-mentioned drives.

There are three phases to the operation of the server: a configuration phase, a boot phase, and an operational phase. During the configuration phase a selected user
 15 with the appropriate administrative pass code can access html pages provided by the server to configure each of the drives in the server. In the example shown an administrator may access the server via a browser on computer 120 for example. After inputting of the appropriate network password the user is provided with HTML pages for configuring the individual drives on the server. The configuration options include:
 20 directory caching, directory and data caching, and archiving. Directory caching provides the ability to browse all discs in a changer without having to cycle through them to read the directories. Not until a file is accessed is there be a need for the changer to load the disc into the drive reader. Directory and data caching is the ability for CD-ROM disks to be mirrored (copied) to a hard disk co-located in a CD-ROM
 25 tower with CD or DVD drives. The data that has been copied is now accessible directly over the network from the hard disk, rather than from the CD itself. The primary advantage of this is access speed. When the user clicks on a file icon, the response back to the user is faster than if the file was being pulled from the CD directly. This is very apparent in the case where the CD drive has stopped spinning or
 30 "idled down" because it has not been used for a period of time. Data access from a CD requires the drive to spin up and then find the file which could take several seconds.

008021 2025260

[illegible]

Figure 6

A

the server. Two web pages 250, 252 are shown. The first of these pages allows user to setup individual CD drives, log events, check network settings, while the second page allows the user to select a specific cache mode for the selected CD/DVD disk or drive. Additional HTML pages (not shown) allow the user to view open and close files during the operational phase of the server. The ability to view files locations, size, and availability without physically accessing any of the associated CD/DVD drives is of particular advantage when the number of drives coupled to the server increases.

FIGS. 3A-F are combined hardware block diagrams of the server and optical storage media shown in FIG. 1 during a write phase of I/O operation. The server portion is shown coupled to a selected CD drive 108 which includes an optical storage media 300 which may be a CD or other drive type. The CD after formatting includes associated media descriptors and format information in header 302. Markers are shown for the first write able address (FWA) "257" and the last write able address (LWA) "1000" for the CD. A plurality of packets 126, 128 from scanner 124 and 136, 138, 140 from scanner 134 are shown entering the network interface 200 of the server portion 106. The following FIGS. 3B-F trace the passage of these packets through the first and second stages 202, 204 respectively of the server to the optical storage media 300.

FIG. 3B shows the lead packets 126, 136 from scanners 124, 134 respectively written into the first stage 202. This stage may be used to accept data packets from multiple users or files or data streams and to coalesce the data from related packets, e.g. those with data from the same data stream or file into larger packets/extents. The processor 206 executing program code stored in memory 208 executes these processes. Additionally, the processor generates the required data structures associated with the selected format chosen for the media to which the data will be written. In the example shown the format for the media is the Universal Disk Format (UDF) which is a subset of a standard for rewritable and write-once media called ISO 13346. This standard enables interchange of data stored on optical media independent of the operating system.

5

20

25

5

10

20

25

The data and control structures previously present in the 2nd stage are flushed to the media prior to the entry of the above mentioned data and control structures.

As is evident from all of FIGS. 3A-F and in particular the above discussed FIG. 3D the input of data, i.e. write requests, are handled serially. Data is initially written to the 1st stage. When that is full it is moved to the 2nd stage. Finally, when the second stage is full it is flushed, i.e. written to the storage media, e.g. CD/DVD. Data is not written over the network directly to either the 2nd stage or the media. By way of contrast data output, i.e. read requests may be handled directly from any one of the 1st stage, the 2nd stage or the storage media. Where a hard drive cache of data is implemented that data may be read directly from the hard drive as well. The processor 206 implements the logic determined by the program code stored in memory 208. This logic effects a the serial movement of data input to the server through the network module 200 through each of the stages to the storage media. The logic implemented by the processor also supports data output from the server through the network module of data to be read from whichever of the stages that data is stored in.

FIG. 3E shows the next phase of the write operation. The extents 138, 128, the ICB 306 and the corresponding VAT 308B have been flushed from the 2nd stage to the optical media over the address range 348-452. The LWA has been updated to the address "453". The last of the extents/packets 140 from scanner 134 has moved from the 1st to the 2nd stage. Since this is the last of the packets associated with the DataStream from scanner 134 the corresponding ICB 304 is also written to the 2nd stage, followed by the VAT 308C.

AXISP001

the management of the 1st and 2nd stages and the optical media will be set forth in the following FIGS. 4AB.

09733707 " 120800

The sequential staged write I/O method and apparatus discussed above in FIGS. 3A-F may be applied to write once or write many media in any of a number of formats without departing from the scope of the invention. It may be applied with and without a hard drive cache. A number of advantages are exhibited by the above discussed I/O apparatus and method: first, multiple users may concurrently read to or write from the storage medium; second, the aggregation and coalescing of data in the first and second stages saves valuable space on the CD. Less of the space associated with lead in and lead out sequences for each block written to the CD is required since block writes are delayed until the 2nd stage is full. Third two data extents belonging to the same file that are sequential can be coalesced in the first stage. This reduces the number of extents for which data structures need to be generated. For communications over the internet the latter feature is of significance due the fact that the typical network packet has payload/data portion limited to 1k. Where a sizeable file, an image file of several meg is being transferred, the ability to coalesce network packets into large extents before writing to the media greatly reduces both time spent writing to and reading from the media and also reduces the complexity of the data structures required to locate the file on the CD/DVD. Fourth, access time to the CD is greatly reduced because fewer write operations to the CD are required.

FIGS. 4A-B are process flow diagrams corresponding with the writing and reading of data to an optical or DVD media as shown in FIGS. 3A-F. Processing begins at start block 400 for which control is passed to decision process 402. In decision process 402 a determination is made as to whether a read request has been received. If so control passes to process 404. In process 404 the location of the data is determined using the VAT and FWA. Control is then passed to process 406. In process 406 the data is read from the location determined in the prior step. Control then returns to decision process 402. Where the next I/O request is not a read request control is passed to decision process 408. In decision process 408 a determination is made as to whether a timeout interval since the last write request has expired. If it has then control passes to process 410 for flushing of the data in the 1st stage to the 2nd

5
10
15
20

25

30

Figure 1. The structure of the proposed model.

corresponding updates are made to the associated ICB. Control is then passed to decision process 440. In decision process 440 a determination is made as to whether any of the flushed extent(s) are the last within the corresponding file/data stream. If so control passes to process 438 in which the corresponding ICB is moved to the 2nd stage as well. Then in process 442 corresponding updates are made to the associated data structures, e.g. ICB and VAT. Control is then passed to decision process 444 which is the same process executed directly after a determination in decision process 440 that the extent being flushed is not the last in the associated file/data stream.

Next in decision process 444 a determination is made as to whether the 2nd stage is full. If not control returns to process 402. If so control passes to process 446. In process 446 the up to date RAM resident copy of the VAT is copied to the 2nd stage. Then in process 448 the 2nd stage is flushed to the media and in process 450 the FWA is updated. Subsequently control returns to decision process 402 for the processing of the next I/O request.

FIG. 4B is a detailed process flow diagram of the processes associated with process block 404-406 discussed above. Processing begins at process block 460 in which the address to be read it is determined. Next in decision process 462 a determination is made as to whether the address to be read is greater than the LWA. If it is control passes to process 464 in which data is read from the 1st stage. Alternately, if the address is not greater than the LWA then in decision process 466 a determination is made as to whether it is less than the FWA. If it is then the requested data resides in and is read from the 2nd stage in process 468. If the results of decision processes 462, 466 are both negative then the data resides on the media 300 (See FIG. 3) and that data is located using the existing file pointers, e.g. VAT and ICB and is provided to the user.

Although the above discussed process flow are in some details specific to the UDF media format the two stage write architecture set forth in the above discussed embodiment may be applied with equal advantage to other media formats without departing from the scope of the invention.

FIGS. 5-7 show various aspects of a user configurable caching implementations that may be practiced on one or more device drivers coupled to the server 106. FIG. 5 is a hardware block diagram of the server, optical and electronic storage mediums shown in FIG. 1. The server 106 includes: a CD driver 212, a hard drive driver 210, the processor 206 and memory 208. Interfacing with the server through the associated drivers is a representative optical storage media device driver 108 with a CD 300 and a hard drive 114 with hard drive media 520. In the example shown the hard drive is formatted with the UDF file structure and includes requisite media identifiers and other directory structures 522. The FWA of the hard drive is "257" and the LWF is address "1000". The CD 300 is formatted with an ISO 9660 data architecture and includes a disk identifier 302 and a FWA of 257 and a LWA of 1000. Other formats may alternately be applied.

The information on the CD includes a plurality of extents 500, 502, 508, 512, and 516. The first of these extents 500 includes the associated ISO 9660 data structures including volume descriptors, directory structures, and the path tables. The second extent 502 includes two data files 504 and 506. The second extent 508 includes a single data file 510. The fourth extent 512 includes the data file 514. The last extent 516 includes a single data file 518.

In the example shown the user has selected directory and data caching for CD 300 on the hard drive media 520. This option is recorded in memory 208 after the user has selected via the above discussed browser interfaces the appropriate cache option. Directory and data caching involves the transfer 550 of the file structure and extent structure to the CD driver. These structures are revised and a directory for the CD is opened on the hard drive. Two files, an extent list 540 and a data file 542 are placed 552 in that directory as a result of that selection. If the user has selected directory caching only then only the extent list and file directories would be copied to the hard drive. The extent list 540 contains offsets for the extents on the hard drive, the data file 542 contains the file structure 530 and the individual files 504, 506, 510, 514 and 518. The space required for the files is considerably reduced due to the removal during the transfer of empty spaces or lead in/out segments which may be as

FIGS. 6A-C show the data structures associated with both the CD and hard drives shown in FIG. 5. FIG. 6A shows the CD 300 with the extents 500, 504 510 514 and 518. Empty space and or lead in and lead out segments separate several of the segments. The mapping of the CD extents E1-5 to extents H1-H4 on the hard drive data file 542 is shown. On the hard drive empty space is removed. Additionally the five segregated extents containing on the CD and the associated files stored in various portions thereof are coalesced to a single file 542 on the hard drive with four offsets. The coalescing of extents speeds up hard drive access, while the removal of empty and or lead-in/out space reduces hard drive size requirements. The use of a single file with offsets greatly speeds access to the hard drive as opposed to multiple files being written thereto.

AXISP001

either the configuration or operational phase of the server. These parameters may be changed at any time to update the cache policy for the associated CD/CD Drive. The selection of cache policy may be implemented using the browser based interfaces shown in FIG. 2.

5 FIG. 6C shows a directory tree on the hard drive 520. A dedicated directory 670 for CD 300 is shown. When directory and data caching are selected for the drive two files 540 and 542 for the extent list and data respectively will be created. When partial caching is implemented both files will be present on the hard drive but the data file 542 will only contain the meta data or file structures 530 (See FIG. 5). This will
10 allow the user to browse the directory and file names for the associated CD from the hard drive. Access to the actual data requires an access to the CD.

 FIGS. 7AB are process flow diagrams corresponding with the boot phase and operational phase of operation of the server shown in FIGS. 5-6AC. Processing and begins at of the start block 700 from which control is passed to decision process 702.
15 In decision process 702 a determination is made as to whether the next CD has been configured for any one of the various cash configurations discussed above. These include: directory caching, directory and data caching and archiving. This decision is made by the server on the basis of cache selections made by the user during system configuration and stored in memory 208 (See FIG. 2). If no cache option has been
20 selected control passes to process 704 for the initialization of the next CD on the SCSI or other bus connected to the server. Alternately, if caching has been enabled control passes to process 706. In process 706 the directory for the selected CD is created on the hard drive. Then in process 708 the list of CD file extents is obtained from the CD data structures. Where the CD is formatted in an ISO 9660 format this information is
25 obtained from the path table. Variations in these processes will apply where alternate formats for the CD are encountered. Next in process 710 additional extents are obtained from the volume recognition descriptors (VRD) on the CD. Then in process 712 the extents which are determined to exist on the CD-ROM and which are contiguous thereon are mapped into a single extent. This step reduces the amount of
30 time required to subsequently access either the CD or the hard drive. In process 714 the extent list 540 (See FIGS. 5, 6B-C) is stored in the directory 670 for the CD 300

on the hard drive (See FIG. 6C). Next, in decision process 716 a CD data file 542 (See FIG. 6C) is created in the above discussed CD directory on the hard drive. The initial offsets for the hard drive extent list fields 662 (See FIG. 6B) are recorded in the extent list. The meta data or file structures 530 (See FIG. 5) for the CD are then
5 copied from the CD into this file. The actual data is not copied at this time.

Next in decision process 718 a determination is made as to whether both directory and data caching has been enabled. If only partial caching is enabled control returns to process 704 for the processing of the next CD drive. Alternately, if directory and data caching has been enabled control passes to decision process 722.
10 The filing processes 722 and 728 result in the calculation of all offset fields 662 by which a file on the CD may be mapped to the hard drive.

In decision process 722 the next extent is processed beginning with process 724. In process 724 data in the extent on the CD is copied to the hard drive at the prior recorded offset. Then control is passed to process 726 in which the next offset
15 for the data in the hard drive is calculated using the prior offset value and the length of the data copied to the hard drive the preceding process. Control is then passed to process 728 in which the calculated offset is registered in the corresponding one of fields 662 in the extent list 540 (See FIG. 6B). Control then returns to decision process 722 for a determination of the existence of the next extent.

FIG. 7B shows the processes associated with read access to the particular CD. Processing begins at start block 750 and proceeds to decision process 752. In decision process 752 a determination is made as to whether a request for the display of a directory of a particular CD has been received by the server. When such a request is received control passes to decision process 754. In decision process 764 a
20 determination is made as to whether directory caching has been enabled in which event control is passed to process 758. In process 758 the extent list and associated directory structures stored on the hard drive as a during partial caching are utilized to rapidly display to the user all of the directories and files structures for the CD. Control subsequently passes to decision process 760. Alternately in decision process 754 when
25 a determination is made that partial caching has not been enabled than control is passed to process 756. In process 756 the disk for the selected CD is spun up and the extent

5

10

25

The many features and advantages of the present invention are apparent from the written description, and thus, it is intended by the appended claims to cover all such features and advantages of the invention. Further, since numerous modifications and changes will readily occur to those skilled in the art, it is not desired to limit the
5 invention to the exact construction and operation as illustrated and described. Hence, all suitable modifications and equivalents may be resorted to as falling within the scope of the invention.

09/33/07 12:08:00